Nulling Interferometry Progress

E. Serabyn
Jet Propulsion Laboratory
4800 Oak Grove Drive
MS 171-113
Pasadena, CA 91109, USA

ABSTRACT

The field of nulling interferometry has seen significant progress over the past several years, in both the conceptual and experimental arenas. Deep, broadband nulling has been demonstrated at optical wavelengths, the techniques have seen initial implementation on telescopes, and the introduction of a symmetric beam-combiner concept has eliminated many of the residual obstacles. Here an overview is provided of promising techniques for effecting the deep cancellation of starlight, and recent results obtained with laboratory and astronomical nulling interferometers are discussed. The next step is the exploitation of nulling techniques at 8-10 m class separated-aperture telescope facilities, and in this vein, a brief overview of the architecture of the Keck Interferometer Nuller is also provided.

Keywords: Nulling, Interferometry, Astronomy

1. INTRODUCTION

Nulling interferometry, a technique which may enable the direct detection of planets around nearby stars, is making great strides forward. New ideas for novel implementations have emerged, important laboratory demonstrations have taken place, initial on-sky observations using single-aperture telescopes have begun, and the installation of facility nullers on state-of-the-art large-aperture interferometers is in the offing. As a result, even though nulling comes with very exacting requirements, and the desired planetary signals are expected to be quite faint, confidence in our ability to employ this technique in future planet-finding observations is growing steadily. In this paper, theoretical ideas for implementing nulling are reviewed, and laboratory and astronomical nullers already deployed or close to deployment are surveyed. As an example of a multi-baseline nulling architecture, the final section touches on the implementation of the first nuller which will use linked 8 – 10 m class telescopes in the near future, the Keck Interferometer Nuller.

2. OVERVIEW OF NULLING TECHNIQUES

In simple terms, the primary goal of nulling interferometry is the accurate subtraction of the electric fields arriving at a number of separate telescope apertures from a distant star[1]. The symmetry and stability requirements which must be met to effect this subtraction are rather exacting[2], including, first and foremost, the introduction of a relative phase shift of π radians between the two beams. Even so, a number of creative approaches to nulling have now been proposed, each with its own advantages and shortcomings. The concepts range from the introduction of achromatic phase retardations by means of dielectric delays, to the introduction of relative field-flips by either geometric or diffractive means, to polarization-based interferometry. This section provides an overview of these approaches. In addition, several approaches which fall short will also be touched upon, in order to elucidate their limitations.

It is instructive to begin with the simplest possible beam-combination approach. Figure 1 shows the light from two telescopes being combined at a single, symmetric, 50/50 beamsplitter[3]. It turns out that there are two fundamental reasons why a deep, achromatic, dual-polarization null cannot be obtained with this straightforward approach. First, because of the high degree of symmetry inherent in this layout, the two beamsplitter outputs O1 and O2 must behave identically at zero optical path difference (OPD), and so at zero OPD each output must receive half the power incident on the two telescopes. Deep cancellation of the incoming light at zero OPD is thus excluded on fundamental grounds, implying that cancellation is possible only at non-zero OPD. Because the best OPD for cancellation of the two light beams is then necessarily wavelength-dependent, deep cancellation over a broad passband is beyond the reach of this simple beam-combiner. Moreover, as is illustrated in Figure 1, the fringes for the two incoming polarization states will

have a relative phase shift of π radians, implying that the two sets of fringes cancel each other completely. This simplest possible beam-combiner is thus necessarily limited to narrow-band single-polarization radiation.

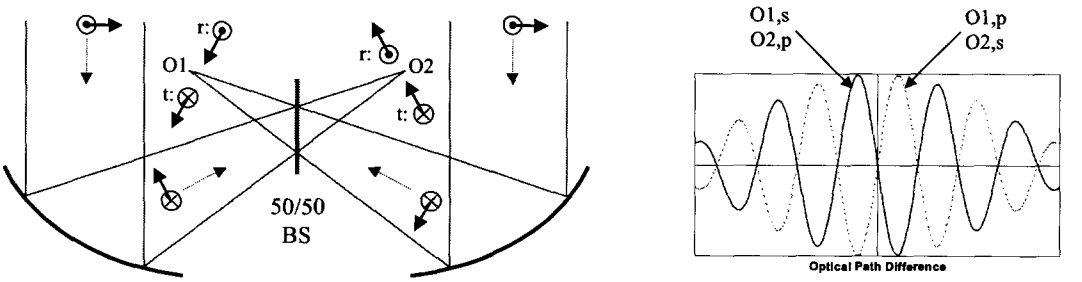

Figure 1. Left: A simple beam-combiner based on an ideal, symmetric, 50/50 beamsplitter. The electric field vectors at the two beamsplitter outputs O1 and O2 are shown excluding the additional $\pi/2$-radian phase shift which is introduced between the reflected and transmitted waves at the beamsplitter. Right: The fringe patterns at O1 and O2. Note that the deepest cancellation does not occur at zero OPD, and that the fringes for the two polarization states (s and p) are out of phase with each other.

As this example illustrates, the details are critical. In particular, the equivalence of the beamsplitter outputs at zero OPD implies that ideal beamsplitters introduce a $\pi/2$-radian phase shift between the reflected and transmitted waves[4], an effect that must be included in the analysis of any candidate nulling beam combiner. Moreover, as is well known, to phase the fringes in both polarization states, it is necessary to break the symmetry of the optical system (in the single-pass beamsplitter case), with an extra mirror reflection in one of the two beam trains (Figure 2). Even so, this can only align the two sets of fringes to first order , since a slight s-p phase shift is introduced by non-ideal reflecting surfaces. As we shall see in Section 2.5, all of these phasing issues can be obviated with a symmetric double-pass beam combiner.

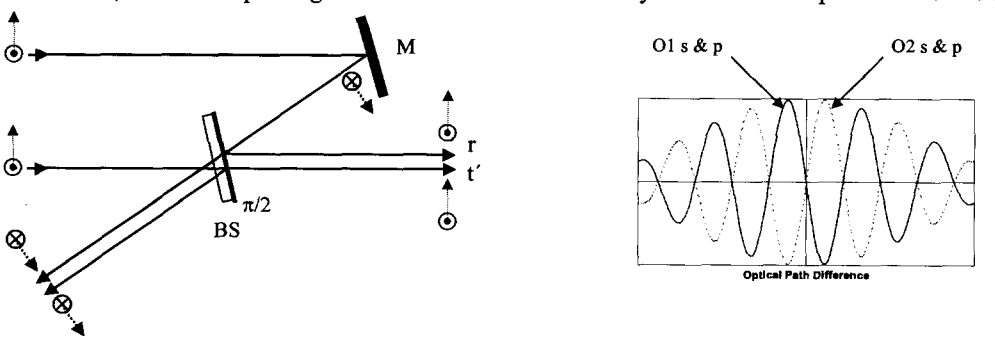

Figure 2. An extra fold mirror (M) in one of the two beam trains phases the fringes in both polarizations. The electric vectors shown exclude the additional $\pi/2$ beamsplitter phase shift (common to both polarizations) between the reflected and transmitted waves.

2.1. Field reversal by phase retardation

Since ideal beamsplitters already insert a relative phase shift of $\pi/2$-radians between the reflected and transmitted beams, in principle only another $\pi/2$ of relative phase is needed to yield reversed fields. One can then seek to produce a deep, achromatic, dual-polarization null fringe by passing the two beams through appropriate dielectrics[5-7], in order to align the first dark fringe at all wavelengths. The simplest approach to introducing the needed dielectric delays in the two-beam case is to use a beamsplitter and compensator with slightly different thicknesses[8] (Figure 3; left). However, this simple case is inherently quite asymmetric, requiring the beamsplitter reflection, **r**, and transmission, **t**, coefficients to have equal magnitudes. In addition, an unbalanced number of both anti-reflection coating traversals and mirror reflections are present, leading to small phase shifts. Deeper nulling thus likely requires a more symmetric layout, as well as multiple glasses in the beam trains, as in the right-hand side of Figure 3. Even so, intensity matching still requires a very symmetric beamsplitter, with $|\mathbf{r}| = |\mathbf{t}|$.

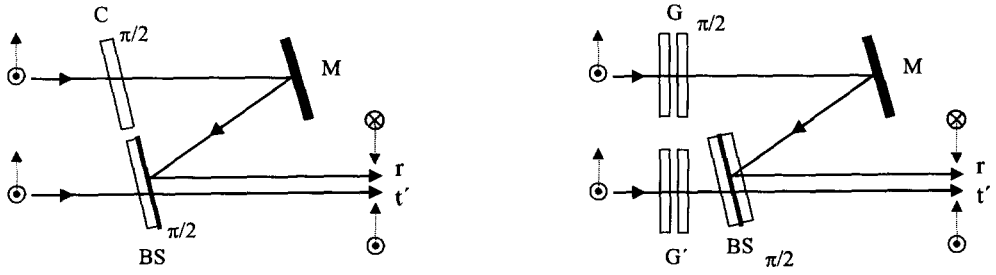

Figure 3. Left: Simple beam-combiner in which the different thicknesses of the beamsplitter, BS, and compensator, C, introduce a relative π/2 phase shift. An additional π/2 of phase is provided by the phase difference between the reflected and transmitted waves at the beamsplitter. Right: A more symmetric version employing a symmetric beamsplitter and multiple glasses. The number of anti-reflection coating traversals is the same in the two beams now, but one extra mirror reflection remains.

An instrument based on the dielectric phase-delay approach has already seen use on a telescope, with atmosphere-limited stellar null depths of order 0.04 attained,[8] and circumstellar dust shells detected in a few cases[9]. The technique has also been used[10] to null a CO_2 laser to a part in 3×10^{-4}. In the case of multiple baselines, it may instead be possible to use vacuum pathlength delays to provide deeper, broader nulls[11].

2.2. Field reversal via geometric field-flips

Alternatively, mirror reflections can be used to geometrically flip the two electric field vectors relative to each other. One implementation of this approach is a rotational shearing interferometer[12-14] (RSI), in which a pair of orthogonal rooftop mirrors flip the fields (Figure 4). Each rooftop mirror reflects the incident field through the line of the rooftop joint, somewhat akin to the action of halfwave plates. The RSI approach can yield a perfect π phase shift, if an extra fold mirror is inserted into each arm to symmetrize the s-plane and p-plane reflections[12,13]. Experiments with RSIs in the JPL nulling lab at both optical and mid-infrared wavelengths have now demonstrated fairly deep nulls (Figure 5), of order 10^{-6} for visible laser light[14], several 10^{-5} for broadband visible light[15], and 10^{-4} for a mid-infrared (9 μm) laser diode[16].

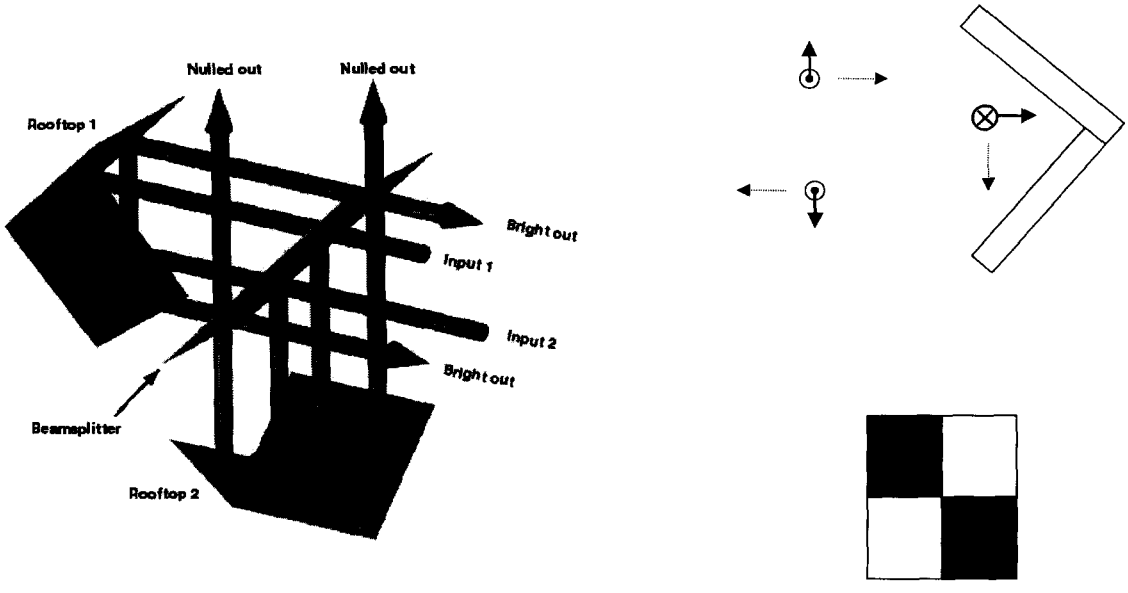

Figure 4. Left: Layout of a crossed-rooftop RSI used as a two-beam nuller. The compensator and two extra fold mirrors have been left out for clarity. Top right: A single rooftop mirror flips that component of the incident electric field vector which is perpendicular to the joint line. Bottom right: A four-quadrant compound beamsplitter which would correct residual asymmetries in the RSI. The beamsplitter coating is on the front side in the two gray quadrants, and on the reverse side in the other two quadrants.

Still, residual asymmetries (unbalanced anti-reflection coating traversals in the beamsplitter/compensator pair, and asymmetric reflections off the two sides of the beamsplitter) are present in this RSI design, which present performance limitations that can only be circumvented with a more symmetric design. For example, one can split the beamsplitter into 4 quadrants, with the pair of beamsplitters on one diagonal reversed relative to the pair on the other diagonal, as indicated in the lower right-hand side of Figure 4. However, this option is rather complex.

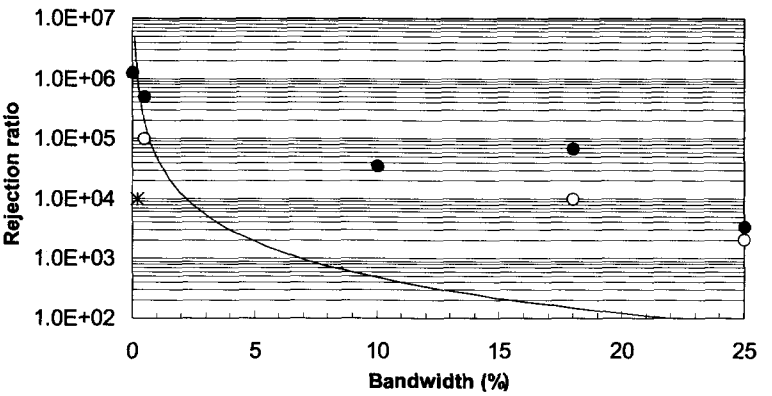

Figure 5. Best measured null depths achieved with the JPL rotational shearing interferometers. The filled circles are the best transient nulls seen at optical wavelengths, and the empty circles are the best stabilized null depths. The asterisk gives the best stabilized mid-infrared null depth achieved. All of the measurements are for linearly polarized light.

Because the beamsplitter is used in double-pass in the RSI approach, the number of output beams is $2^2 = 4$. This doubling of the outputs is not necessarily a disadvantage, as unity efficiency can be achieved by summing over the two nulling outputs, and the extra outputs may in fact be put to other uses (control loops, operation at other wavelengths, etc.) More importantly, in contrast to the phase-retardation approach, the field-of-view is inverted along with the fields in an RSI, so that off-axis sources are doubled in the combined beam, a less than optimal situation.

An alternative starting point is to consider introducing a relative field flip into a pair of beams before a beamcombiner is encountered, such as with a pair of inverted, right-angle, reflective periscopes[17] (Figure 6). Orthogonal half-wave-plates of a variety of types can also be considered (e.g., doubled Fresnel rhombs, birefringent crystals, and form-birefringence-based waveplates), but achromaticity is an issue, and experimental nulling work with such devices has been limited. Of course, the introduction of a relative field flip before the actual beam combiner then necessitates the use of a constructive beam combiner, in order to avoid the superfluous introduction of additional phase shifts. Therefore, a different optical beam combiner layout is also called for, the discussion of which is deferred to Section 2.5.

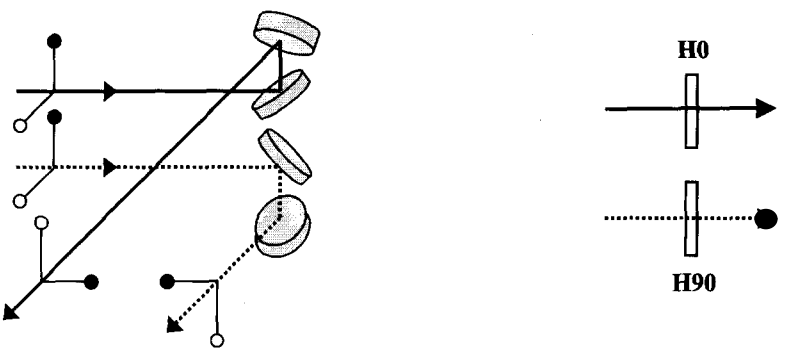

Figure 6. Left: An inverted pair of right angle periscopes that provide a relative field flip. Note that each arm has one s-plane and one p-plane reflection, so that the system is symmetric with respect to polarization state. Right: A pair of orthogonal achromatic half-wave plates will have a similar effect.

2.3. Field reversal via diffraction

Another method of reversing electric field vectors relies on passing one of the two beams to be combined through an extra focus[18,19], since passage through focus introduces an achromatic π phase shift into the beam[20]. This field reversal can be understood in a heuristic way as the result of the pupil inversion which occurs upon passage through focus. More rigorously, this phase shift arises in the $\pi/2$ phase which appears explicitly in the Huygens-Fresnel and Kirchhoff diffraction integrals[20]. In the case of a passage through focus, one factor of $\pi/2$ is introduced upon taking the beam from collimated space to the far-field (i.e., to the focal plane), and a second factor of $\pi/2$ is impressed upon the beam upon its return to collimated space. Thus, although this field reversal superficially resembles a geometric field flip, it is in fact induced by diffraction.

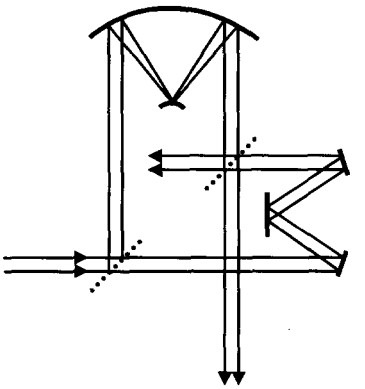

Figure 7. Layout for a nuller in which one beam acquires an achromatic phase shift of π radians upon passing through an extra focus.

One proposed implementation of this approach[18,19] (Figure 7) involves replacing the two end mirrors in a Michelson interferometer with cat's eye reflectors. The beam in one arm passes through a focus, but not in the other. The obvious asymmetry (reflections at mirrors of different radii of curvature) is a limitation at small focal ratios, but not at large enough focal ratios. Other than that, since this approach also relies on a field flip, the issues are much the same as for an RSI. Such an instrument has also already been operated at a telescope[19], with stellar nulls at a wavelength of 2 μm of order 0.02 achieved. More importantly, the detection of a companion within the central Airy spot of a telescope's diffraction pattern has been demonstrated[19], much closer to the axis than a classical coronagraph would allow.

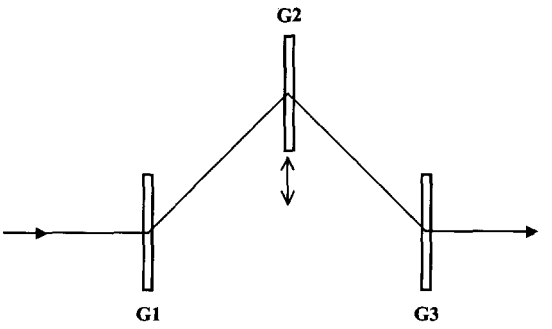

Figure 8. An achromatic grating-based phase shifter. The output phase depends achromatically on the vertical position of G2.

Another diffraction-based approach to achromatic phase shifting relies on translating a transmission grating laterally (Figure 8). As the chief ray moves across the grating facets, the output phase cycles through 2π in achromatic fashion[21]. A pair of these units would thus allow for the introduction of an achromatic relative field flip. This approach has not yet been applied to nulling interferometry, likely because of efficiency and alignment issues.

2.4. Polarization interferometry and Pancharatnam phase

Perhaps the simplest method of differencing light beams is with linear polarizers[22]. Parallel linear polarizers can be used to select a common polarization state at a pair of telescopes, and then this polarization component can be rotated into opposition in the two beam trains with additional polarizers (Figure 9; left side). The cancellation should be achromatic since only field orientation is involved, but this approach is very lossy – summing both beamsplitter outputs, only 1/8 of the light survives. Replacing the second and third polarizers in each arm with half-wave plates would increase the efficiency, but since the first pair of polarizers would then also no longer be needed, we are back to the (unity efficiency) crossed half-wave plate case (Figure 6). The linear polarizer case thus need not be considered any further.

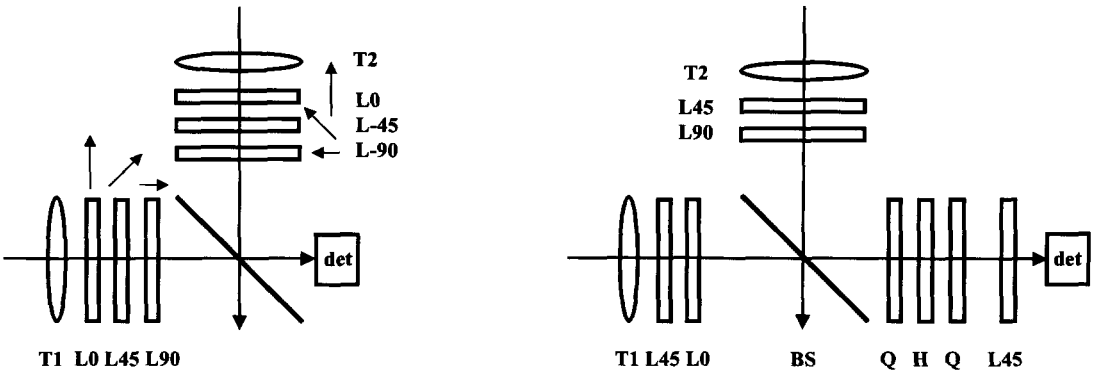

Figure 9. Polarization interferometry. Left: First a common linear polarization state is selected by linear polarizers L0 at telescopes T1 and T2. This polarization component is then projected in opposite directions in two ±45° increments by successive linear polarizers, leaving the fields opposed. Right: A single common linear polarization state at 45° is first selected at telescopes T1 and T2. This polarization component is then projected once in opposite directions by ±45° to produce orthogonal linear polarizations, before passing through a quarter-wave plate, Q, a half-wave plate, H, a second half-wave plate, Q, and a final linear polarizer, L45, all oriented at 45°. Neglecting the beamsplitter phase, the first Q would produce oppositely oriented circular polarization states. H then reverses these states, the second Q restores the original linear polarization states, and the final linear polarizer projects these back onto a common axis for interference. Rotation of the HWP by ±45° from nominal introduces a relative phase shift of π radians.

Alternatively[22,23] (Figure 9; right side), after selecting a common linear polarization state at a pair of telescopes, the polarization vectors can be taken around different circuits on the Poincaré sphere. If the circuits enclose finite solid angles, a pathlength-independent Pancharatnam phase[24], analogous to the Berry phase in quantum mechanics[25], is introduced. With properly oriented polarizers and wave plates (Figure 9) the desired field reversal can be effected (if the beamsplitter phase is neglected in this layout). It is in fact not necessary to invoke the Poincaré sphere formalism to understand the result, as tracking the polarization vectors through the successive devices algebraically is also quite straightforward. This case again has low (1/8) efficiency, although factors of 2 can be recovered both by replacing the second pair of linear polarizers with half-wave plates, and by duplicating the assembly for the orthogonal polarization component. An undesirable complication in the case shown in Figure 9 is that the selected field component arrives at the beamsplitter from the two arms with orthogonal polarizations, leaving the beamsplitter phase and amplitude matching issues fairly intractable. Once again, a double-pass beamsplitter configuration would obviate these problems (Section 2.5). Experiments verifying aspects of the Pancharatnam phase have been performed with a single beam of light[22], but an experiment analogous to the two-beam astronomical beam combiner case has yet to be carried out.

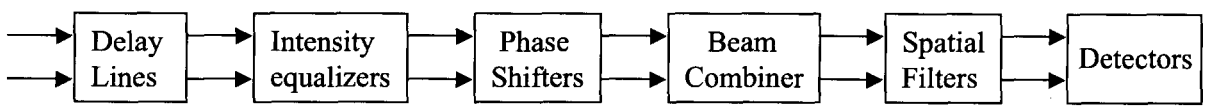

Figure 10. Functional block diagram of an ideal nulling beamtrain.

2.5. The fully symmetric nuller

To overcome the asymmetries and limitations inherent in the various nuller implementations described above, a more symmetric approach is needed. Also advantageous would be the separation of the field-flip and beam combination stages into sequential steps, and the ability to null both polarization components simultaneously. Ideally one would have complete separation of the optical functions into a sequential block diagram as in Figure 10.

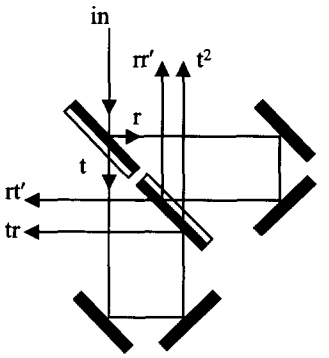

Figure 11. Reciprocal beamsplitter arrangement allowing the production of the equivalent **rt** and **rt'** outputs.

To achieve complete symmetry, the two beams to be combined need to be treated equivalently by the beamcombiner. This is incompatible with a single-pass beamsplitter for several reasons: the requisite extra reflection in one arm (Figure 2), the beamsplitter's $\pi/2$ phase shift between **r** and **t**, and the requirement that $|\mathbf{r}| = |\mathbf{t}|$. On the other hand, with a reversed double-pass beamsplitter arrangement, as e.g. in Figure 11, a completely symmetric treatment of the two beams can be achieved in quite simple fashion[17], since the product of the beamsplitter reflection and transmission coefficients **rt** and **rt'** are in fact identical. In passing through a pair of reversed beamsplitters, each beam thus acquires identical phase shifts, thereby completely removing the beamsplitter phase issue from the problem. With equal phase shifts, the two beams interfere constructively at zero OPD, so that a reversed double-pass beamsplitter arrangement yields a constructive beam combiner. There are in fact three possible variants of symmetric, constructive beam combiners[17], based on modified Michelson, Sagnac, and Mach-Zehnder interferometers (Figure 12). All rely entirely on flat optics.

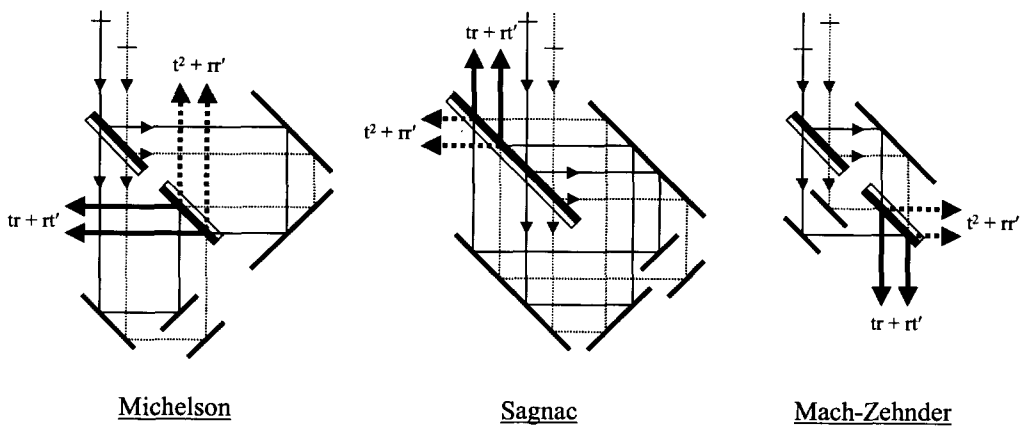

Figure 12. Three fully-symmetric constructive beam-combiner layouts.

The second necessary ingredient is a perfectly symmetric and achromatic method of reversing the fields prior to the beam combiner. Two example have already been shown in Figure 6: a mirror-symmetric pair of right-angle periscopes, and general orthogonal, achromatic half-wave plates. In either case, the field reversal has the effect of converting the symmetric **rt** + **rt'** outputs into asymmetric **rt** - **rt'** outputs, which are identically zero for lossless beamsplitters. Combining these two facets of the solution led to the concept of the fully symmetric nuller[17].

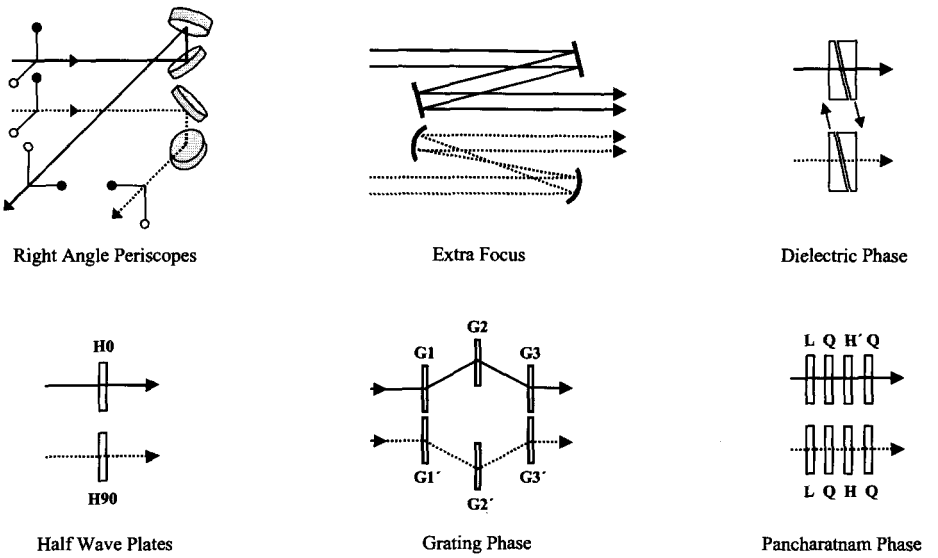

Figure 13. Six types of phase shifters that can be inserted prior to a symmetric, constructive beam-combiner to convert it into a nuller. The devices based on right-angle periscopes and an extra focus invert the image, and so can only provide a phase shift of π radians. The half-wave plates also provide a phase shift of π radians, but do not invert the image. The remaining three examples can provide arbitrary phase shifts, and do not invert the image field. The Pancharatnam phase shifter is the only one of these which is inherently single-polarization. A relative rotation of 45° between H and H′ provides the needed π phase shift in the Pancharatnam phase shifter.

Note that with a reciprocal pair of beamsplitters, the actual beamsplitter performance is fairly unimportant, as long as both beamsplitters are part of a matched set, with the matching criteria applying both to the coatings and the substrates (thicknesses, homogeneities, etc.) Thus, off-the shelf coating designs become possible. In principle, because each recombining beam sees identical reflections and transmissions, the resultant perfect subtraction is then independent of wavelength, polarization state, angle of incidence, etc. This approach thus has a sizable number of advantages. Of course real world issues, such as substrate wedge angles and coating non-uniformities, can re-introduce slight asymmetries.

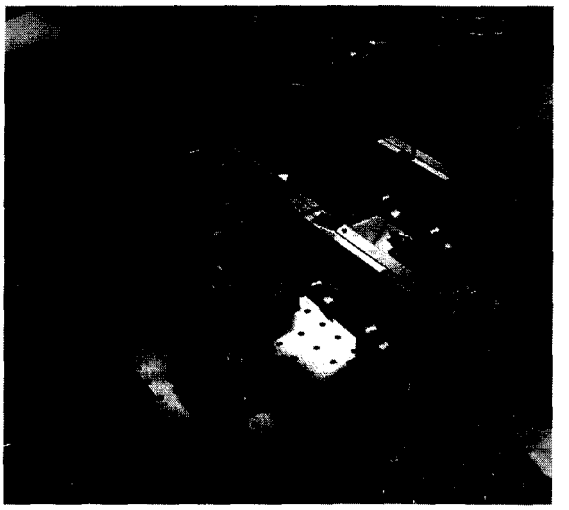

Figure 14. An assembled modified Mach-Zehnder beam combiner at JPL. Four rather than two beamsplitters can be seen in a row near the center of the image because each of the two reversed beamsplitters has been split in two to reduce cost.

In fact, it is easy to generalize the field-flip/constructive-combiner concept to allow for other field-reversal approaches, in particular to those which depart slightly from perfect symmetry. Dielectric phase delays provide a case in point – a pair of phase plates, symmetric except for their achromatic π phase difference (Figure 13; top right), when followed by a symmetric beam combiner, provides an arrangement which cures all of the asymmetry-related problems discussed in Section 2.1, and successfully separates the phase plate and beamsplitter design issues, thereby greatly easing both (e.g., the $|\mathbf{r}| = |\mathbf{t}|$ beamsplitter requirement is gone). In fact, any of the six field-reversal approaches shown in Figure 13 (right-periscopes, an extra focus, dielectric delays, crossed half-wave plates, grating phases, and Pancharatnam phase) can be utilized prior to a symmetric beamcombiner. Indeed, the advantages of this concept are so great that the idea of a near-achromatic field reversal prior to a symmetric beam combiner has now achieved wide acceptance by many of the parties involved in planning future nulling experiments.

2.6. Performance verification of the symmetric nuller concept

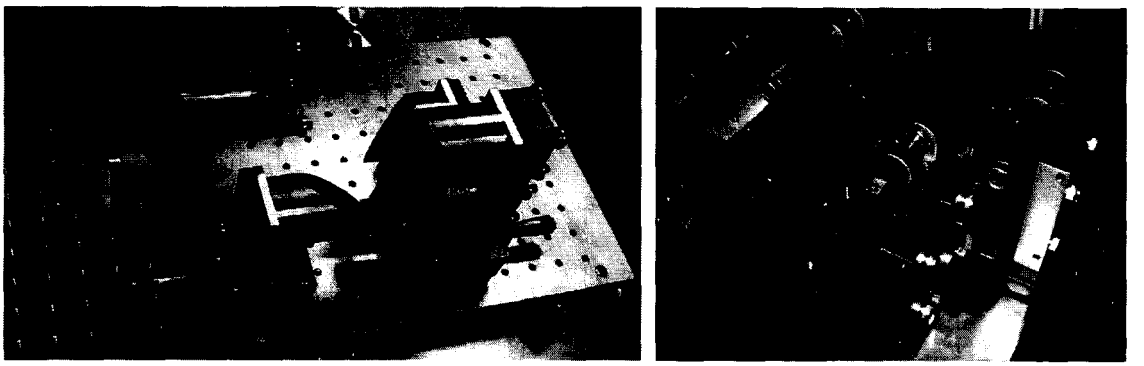

Figure 15. Two field-reversal units built at JPL. Left: Assembly of a pair of mirror-symmetric right angle periscopes. Right: Two sets of dielectric phase-delay plates[26], as in Figure 13, made of ZnSe glass.

Once the concept of the fully symmetric nuller was introduced, we set about building such devices in the JPL nulling lab. Our two coupled goals were the delivery of an operational nuller to the Keck Interferometer, and technology development for the Terrestrial Planet Finder mission. As both of these experiments will operate at mid-infrared wavelengths ($\approx$ 10 μm), all of our experiments to date with fully symmetric nullers have been in the mid-infrared. We selected the modified Mach-Zehnder beamcombiner because it is the most compact, with the fewest mirrors (Figure 12). One of our assembled modified Mach-Zehnder interferometers is shown in Figure 14. Our experiments have employed both periscopic field-flippers and dielectric delays (Figure 15). Indeed, since at the Keck Observatory dielectric phase correctors will be required as a matter of course to correct for atmospheric fluctuations expected to be well above π radians[26], it quickly became clear that a field flip (a mere π radians) would only be a small part of the total phase correction needed on site, implying that a system capable of a perfect field-flip is largely superfluous when observing through the atmosphere. Our experimental work has thus recently concentrated on the dielectric phase-delay approach.

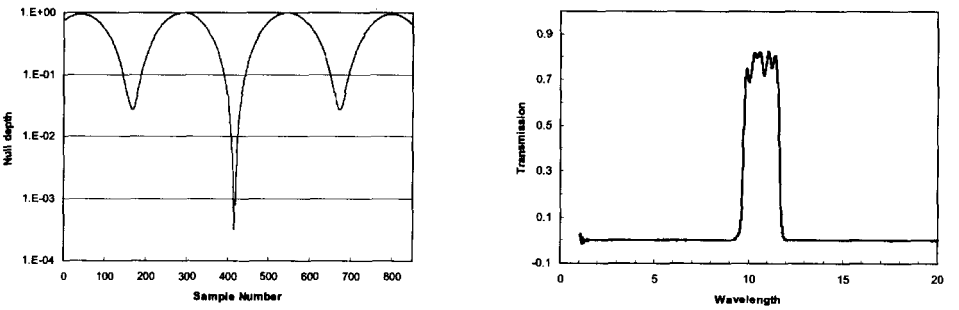

Figure 16. Left: Scan in optical path through the null fringe for dual-polarization mid-infrared thermal radiation in one of the JPL modified Mach-Zehnder nullers. The field reversal was provided by dielectric plates tuned to symmetrize the fringes to either side of the deep central null fringe. Right: The filter passband. The passband center is at 10.7 μm and the FWHM is 18%.

Although working at mid-infrared wavelengths makes the experimentation more challenging than at optical wavelengths, good progress has in fact been possible, and mid-infrared laser-diode nulls of 3×10^{-5} have already been obtained. More significantly, we have in fact been able to verify the essentials of the approach, since nulls better than 8×10^{-4} have been achieved for broadband thermal radiation from a filament (Figure 16). The thermal radiation was filtered to a bandwidth of 18%, centered at a wavelength of 10.7 μm. No polarizers were needed to achieve these rejection levels, thus verifying the polarization-independence of the symmetric approach. Our best mid-infrared laboratory results to date obtained with a modified Mach-Zehnder system are summarized in Figure 17.

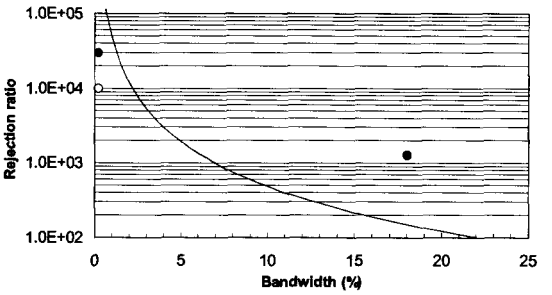

Figure 17. Best measured mid-infrared nulls to date. The laser diode operates at 9.2 μm, and the 18% passband for the thermal radiation is shown in Figure 16. The broadband result is for dual-polarization light. The curve shows the best rejection possible at the first dark fringe for a symmetric, constructive combiner. Our result for broadband radiation is thus well into the achromatic regime.

3. ANCILLARY ISSUES

The beam combiner is of course only one part of the system (Figure 10), and several ancillary issues must be faced by a real system in which two (or more) beams traverse different optical paths. All of these cannot be explored here, but a few examples will be given to suggest that solutions to many issues are not as difficult as commonly perceived.

3.1. Symmetric beam injection

To test a nuller's capabilities, two very symmetric input beams must be available. Fortunately, generating the requisite beams is very easy with dual-beamsplitter designs – one can simply build a second nuller and operate it in reverse to split the light perfectly symmetrically. Even better, one can simply reverse-illuminate the nuller in question with a single beam sent into one of its nulling outputs, and allow the nuller beamsplitters themselves to produce two fully symmetric output beams exiting the nuller in reverse (Figure 18). Since the two resultant identical beams exit through the desired input ports, they can then be retro-reflected to pass through the nuller in the forward direction. In the second pass, two beams enter the nuller through the normal input ports, and can then interfere as in normal operation. This is in fact how the deep mid-infrared nulls discussed above were obtained. This reverse-injection scheme does not work with the field-flip approach, as the fields would be flipped twice in double pass. This is another reason for turning to dielectric phase plates, or some other approach for which the phase-shift can easily be cut in half.

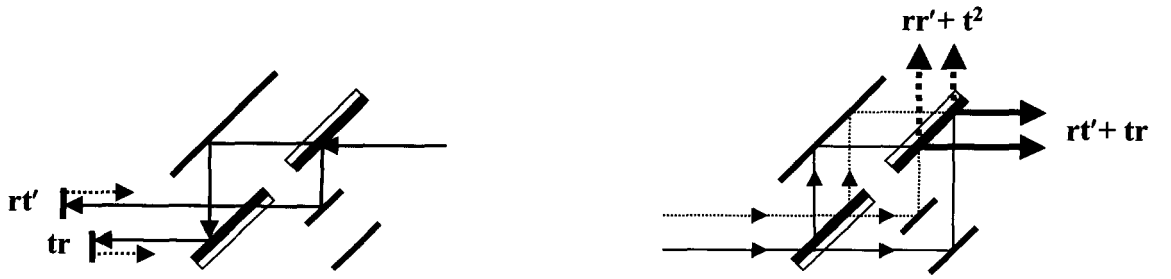

Figure 18. Scheme for producing two beams of suitable symmetry for nulling. Left: The nuller is reverse-illuminated by a single input beam. Right: the two output beams are retro-reflected to pass through the nuller again, this time with interference taking place.

3.2. Intensity balancing

Another area of importance is intensity matching. Even though this must be done achromatically, this is not as difficult as generally perceived, at least for nulls of order 10^{-4}, for which the intensities need only match to a few percent. Two very simple approaches can be considered: changing the relative pupil diameters before recombination, and vignetting a tiny part of the aperture. Of these, the second is the more achromatic because it simply scatters a given geometric fraction of the power out of the beam, while the first changes the diameter of the point spread function, which is a chromatic quantity. To vignette the aperture, two very simple schemes come to mind (Figure 19): a narrow "scissors-like" obscuration across the center of the aperture, which can widen or narrow akin to a pair of scissors, or a small rotating vane across the center of the beam, which resembles one slat of a "Venetian blind". The second is obviously easier to implement, and we have in fact already implemented a very simple version of a Venetian-blind intensity modulator - a thin metal ruler taped vertically to a block of metal.

Figure 19. Two variants of intensity modulators. Left: "scissors." Right: "Venetian blind".

3.3. Spatial filtering

While spatial filtering is generally needed to achieve deep nulls, the fact that wavefront quality increases as the wavelength increases implies that the need for spatial filtering eases significantly in the mid-infrared. It should thus be possible to attain fairly deep nulls in the mid-infrared without heroic polishing efforts and without the use of spatial filtering. If wavefront quality dominates the nulling error budget, and the null depth is given by $n\ (\pi\ \sigma_{HeNe})^2\ (\lambda_{HeNe}/\lambda)^2$, where n is the number of reflections, σ_{HeNe} is the individual surface rms in units of a HeNe wavelength, λ_{HeNe} is 632 nm, and λ is the observing wavelength, then even standard tenth-wave quality (peak-to-valley at 632 nm) optics can provide a null depth of order 10^{-4}, assuming 10 reflections. Nulling imaging, in which no spatial filter is employed, thus has a good deal of potential at long wavelengths.

3.4. Null stabilization

Usable nulling approaches must also provide a robust method of locating and stabilizing the null. Two methods of null stabilization have already been demonstrated at optical wavelengths, pathlength dithering[15] and quadrature output stabilization[27]. The first of these techniques was demonstrated using only a faint white light signal comparable to the stellar flux incident on 30 cm apertures in the optical. In fact, in the signal dominated regime, the sensitivities of both approaches are comparable[27]. The primary approach under consideration for mid-infrared wavelengths, where background fluctuations dominate the noise, is based on secondary (shorter) wavelength sensing. The wavelength split can occur either before the delay lines (Section 4.1), or, if chromatic phase-shifts are introduced, at the short-wavelength quadrature outputs of the nuller[10]. Note that all of these stabilization approaches operate either at an extremum or an inflection point of the fringe. Algorithms for stabilizing at arbitrary phases[11] remain to be defined, and likely must rely on metrology. Of course, metrology is likely to be incorporated into all stabilization approaches.

4. GROUND-BASED MULTI-BASELINE NULLING

The discussion thus far has concentrated largely on the physics of the nulling process, implicitly assuming a single-nulling baseline. Multi-baseline nullers can be used to broaden either the spatial or spectral character of the null[5,11,28,29], or to effect a double subtraction. The former is likely applicable to the space-based case, while the latter is particularly important for the ground-based case, where both the star and the thermal background need to be suppressed. In principle, the sky background can be suppressed either by implementing synchronous on/off spatial chopping, or by means of interferometric pathlength modulation. In the former case, the pathlength control loops would be broken at the chopper frequency, while in the latter case the control loops can operate continuously. The Large Binocular Telescope[9] (LBT) is slated to use spatial chopping, while the Keck Interferometer Nuller will rely on interferometric modulation.

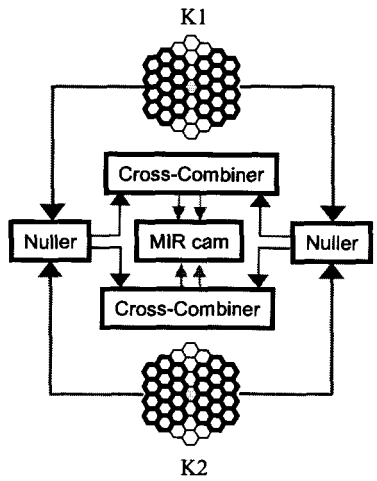

Figure 20. Architecture of the Keck Interferometer Nuller. The primary apertures are split into two (highlighted) subapertures each, and the star is nulled on the two long, parallel baselines between the corresponding subapertures. The residual signal from e.g. exozodiacal light is then detected by a mid-infrared camera by scanning the cross-combiners over a wavelength.

4.1. The Keck Interferometer Nuller

The Keck Interferometer Nuller (KIN) is based on sub-aperturing the telescopes to create a pair of long, parallel nulling baselines (Figure 20) between corresponding subapertures. After the star is nulled on each of these long (85 m) baselines, the outputs of the pair of nullers are fed to a pair of "cross-combiner" beamsplitters, which constitute a pair of standard astronomical interferometric beam combiners. The cross-combiners effectively combine the light on the short baselines across the Keck primary apertures, and operate in standard fashion by repeatedly scanning the post-nuller OPD over a distance of a wavelength to measure the non-stellar source flux. Thus, the data product is a standard fringe scan.

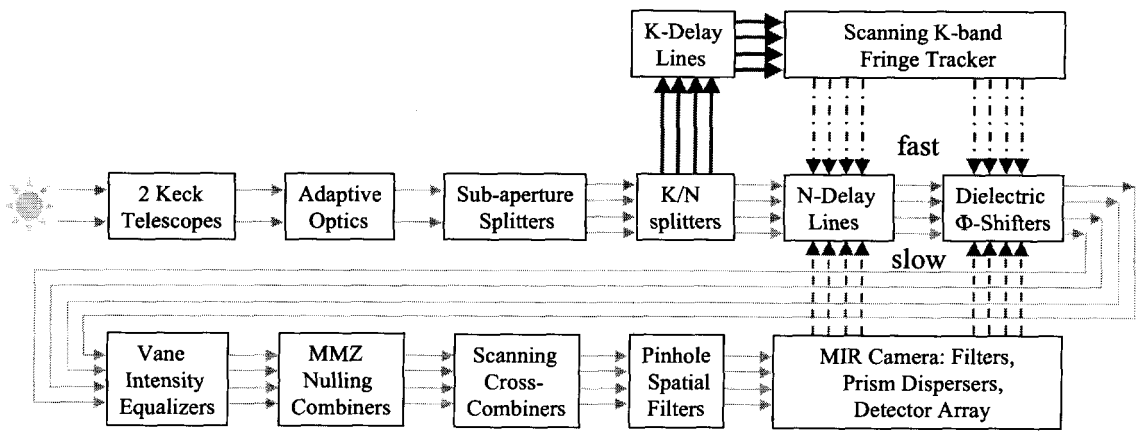

Figure 21. Functional block diagram of the Keck Interferometer Nuller. The solid arrows are light signals (heavy arrows for K-band light after the K/N splitter), the dash-dotted arrows are (fast) control signals from the K-band (2 μm) fringe tracker to the N-band (10 μm) delay lines and phase shifters, and the dashed arrows are (slow) control signals from the mid-infrared camera to the same.

A more detailed functional block diagram of the KIN is provided in Figure 21. Most of the components are self-explanatory, but the control scheme bears further elaboration. Rapid fringe tracking (scanning in pathlength by a wavelength) is done with K-band (2 μm) light in a set of dedicated delay lines, and the resulting fringe position information is then fed forward to the separate delay lines and dispersion compensation system used by the N-band (10 μm) light, which are targeted to be held at null.

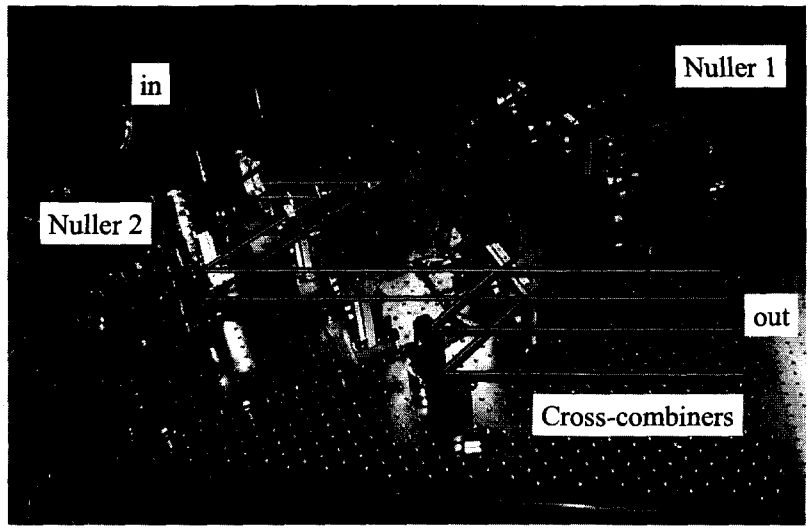

Figure 22. Photograph of the Keck Interferometer Nuller optics plate. The nullers and crosscombiners are labeled. The four beams from the four Keck subapertures reach the nulling breadboard at the top left. After nulling and cross-combination, four beams depart at the lower right to the mid-infrared camera.

The actual KIN nulling/cross-combiner breadboard is now assembled (Figure 22) and is undergoing actuation and testing, after which it will be integrated with the other system components (the mid-infrared camera, delay lines, and K-band fringe tracker), and then shipped to the Keck Observatory next year.

4.2. Science goals

The 85 m length of the Keck baseline implies that for solar type stars at a distance of about 10 pc, the null depths attainable with the KIN in the 10 μm atmospheric window will be limited to $\approx 10^{-3}$ (Figure 23; top left). One of the primary goals of the KIN is the detection of exozodiacal emission from nearby stars; this has already been explored in detail elsewhere[30], and the calculations indicate a sensitivity limit for the KIN per night of integration of several times the brightness of our own zodiacal cloud (our own zodiacal cloud is roughly 10^{-4} times as bright as the sun in the mid-infrared). Due to the KIN's narrow fringe spacing (24 mas at 10 μm; Figure 23, bottom left), and the steep radial temperature dependence of any exozodiacal dust present, the KIN's predominant sensitivity will be to emission arising very close to the star[30] – specifically at about a tenth of an AU for a star at 10 pc (Figure 23; top right).

Another area in which the KIN has unique potential is hot exo-planet detection. In particular, because the baseline length puts the first constructive fringe so close to the star (Figure 23; bottom left), the first transmission maximum actually allows good access to the brightest category of exoplanets: close-in Hot Jupiters. Indeed, at maximum elongation, several known[31] Hot Jupiters lie close to the KIN's first fringe maximum at 5 μm, and at 10 μm the transmission is only down to about 0.2. A search for such objects is made easier by their relatively favorable contrast ratios (the blackbody ratio is $\approx 10^{-3}$ for 1200 K bodies of Jupiter's size; Figure 23, bottom right), and also by their short orbital periods, which can provide a day-to-day on/off signal modulation.

One very important question which then can be addressed with the KIN is the excess (above a single-temperature blackbody) of emission predicted[32] to be present in planetary atmospheres in the vicinity of 5 μm (due to high clear-

atmosphere transparency at those wavelengths). Large excesses should be easily detectable if the pathlengths can be held stable enough for such short wavelength observations, but even upper limits at the level of several parts per thousand can provide useful constraints on any possible excess emission near 5 μm, and so on exo-planet atmospheric transparencies, i.e., on the atmospheres' cloudiness and dustiness. Currently this is not part of the baseline KIN plan, but it is clearly an area with potential.

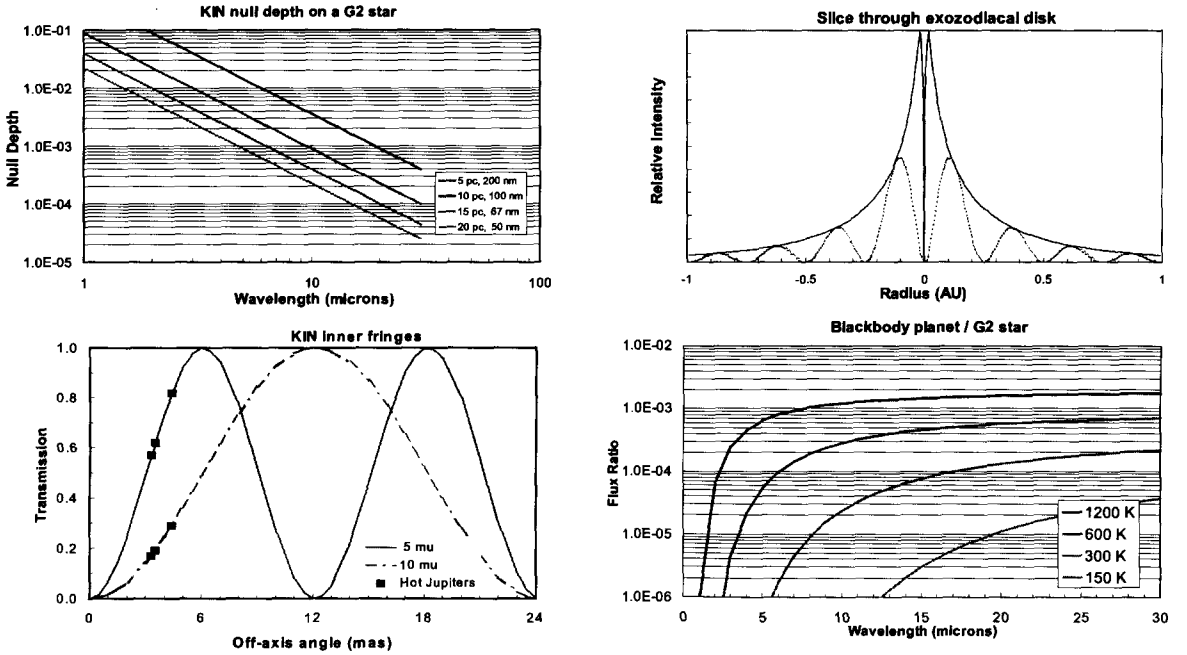

Figure 23. Top left: Null depths vs. wavelength and stellar distance for the KIN. Also given is the equivalent OPD stability needed. Top right: A cut through a zodiacal cloud model (solid) and the product of this model with the KIN fringe transmission (dotted) for a star at 10 pc. Bottom left: The KIN's innermost fringes at 5 and 10 μm. The data points locate the major axes of the orbits of several Hot Jupiters (left to right: τ Boo, 51 Peg, υ And b). Bottom right: Wavelength-dependence of the contrast ratio of blackbody exoplanets (with Jupiter-sized diameters, but various temperatures) to solar-type stars.

5. CONCLUSIONS

Great progress is occurring in the field of nulling interferometry, in the experimental, observational, and conceptual arenas. Significantly, a new nulling architecture, in which a field reversal is introduced prior to a symmetric, constructive beamcombiner, has emerged. One result is that nulling instruments are beginning to bear more of a resemblance to standard astronomical interferometers than hitherto. All of these ideas and experimental results have allowed a greater confidence in our ability to exploit this high-precision technique as we begin to apply it on the largest class of linked telescopes currently (and soon to be) available. Interestingly, a common nulling architecture – dielectric field-flip stages followed by a symmetric modified Mach-Zehnder beam combiner – seems to be gaining wide acceptance as the default nulling architecture for future experiments and missions.

ACKNOWLEDGEMENTS

I would like to thank the members of the JPL nulling team, and in particular Sam Crawford, Michelle Creech-Eakman, Jason Geis, Chris Koresko, Stefan Martin, and Bertrand Mennesson for their hard work in obtaining the various experimental results discussed herein, and in helping to move the concepts forward. Special thanks are also due to Mark Colavita and Stefan Martin for unveiling the mysteries of Microsoft Word to me. I also thank the Keck Interferometer and Terrestrial Planet Finder projects for generous support. This work was carried out at the Jet Propulsion Laboratory, under contract with the National Aeronautics and Space Administration.

REFERENCES

1. R.N. Bracewell, Nature 274, 780 1978
2. E. Serabyn, Proc. SPIE 4006, 328, 2000
3. R.N. Bracewell & R.H. MacPhie, Icarus 38, 136, 1979
4. W.A. Traub, in "Principles of Long Baseline Stellar Interferometry", JPL Pub. 03-01, 31, 2001
5. R. Angel, in "The Next Generation Space Telescope", 81, STScI, 1989
6. R.M. Morgan, J. Burge, & N. Woolf, Proc. SPIE 4006, 340, 2000
7. A.L. Mieremet, J. Braat, H. Bokhove & K. Ravel, Proc. SPIE 4006, 1035, 2000
8. P.M. Hinz et al., Nature 395, 251, 1998
9. P.M. Hinz, these proc.
10. P.M. Hinz, R. Angel, N. Woolf, B. Hoffmann & D. McCarthy, Proc. SPIE 4006, 349, 2000
11. A.L. Mieremet & J.J.M. Braat, Applied Optics 41, 4697
12. M. Shao & M.M. Colavita, ARAA 30, 457, 1992
13. E. Serabyn, J.K. Wallace, G.J. Hardy, E.G.H. Schmidtlin & H.T. Nguyen, Applied Optics 38, 7128, 1999
14. S.R. Martin , E. Serabyn & G.J. Hardy, these proc.
15. K. Wallace, G. Hardy & E. Serabyn, Nature 406, 700, 2000.
16. C. Koresko, B. Mennesson, S. Crawford, M. Creech-Eakman, K. Wallace & E. Serabyn, these proc.
17. E. Serabyn & M.M. Colavita, Applied Optics 40, 1668, 2001
18. P. Baudoz, Y. Rabbia & J. Gay, A&A Suppl. 141, 319, 2000
19. P. Baudoz et al., A&A Suppl. 145, 341, 2000
20. M. Born & E. Wolf, Principles of Optics, 6th ed., Sect. 8.8.4, Pergamon, 1984
21. N. George & T. Stone, Optics Communications 67, 185, 1988
22. N. Baba, N. Murakami & T. Ishigaki, Optics Letters 26, 1167, 2001
23. W.J. Tango & J. Davis, Applied Optics 35, 621, 1996
24. S. Pancharatnam, Proc. Indian Academy of Sciences A 44, 247, 1956
25. M.V. Berry, J. Modern Optics 34, 1401, 1987
26. C. Koresko, B. Mennesson, E. Serabyn, M. Colavita, R. Akeson & M. Swain, these proc.
27. E. Serabyn, S.R. Martin & G.J. Hardy, Proc. 2001 IEEE Aerospace Conference, 4-2027, 2001
28. J.R.P. Angel & N.J. Woolf, ApJ 475, 373, 1997.
29. B. Mennesson & J.M. Mariotti, Icarus 128, 202, 1997
30. E. Serabyn, M.M. Colavita & C.A. Beichman, in "Thermal Emission Spectroscopy and Analysis of Dust, Disks and Regoliths", ASP Conf. Ser. Vol. 196, 357, 2000
31. http://exoplanets.org
32. A. Burrows et al., ApJ 491, 856, 1997